IDA

INSTITUTE FOR DEFENSE ANALYSES

High Performance Computing Modernization Program: Program Review–PR 2000

Alfred E. Brenner, Task Leader

June 2000

IDA Document D-2486
Log: H 00-001965

This work was conducted under contract DASW01 98 C 0067, Task AK-5-1877, for the Office of the Deputy Under Secretary of Defense (Science and Technology). The publication of this IDA document does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that Agency.

INSTITUTE FOR DEFENSE ANALYSES

High Performance Computing Modernization Program: Program Review–PR 2000

Alfred E. Brenner, Task Leader

Preface

This document was prepared by the Institute for Defense Analyses under the task order, High Performance Computing Modernization Program (HPCMP) Review, to assist in the development and execution of a periodic program review process to evaluate the performance of the HPCMP, to participate in the FY2000 program review, and to document and analyze the conclusions of the review. This report was prepared for the Director, High Performance Computing Modernization Office, Office of the Deputy Under Secretary of Defense (Science and Technology) (DUSD (S&T)).

This document was reviewed by Dr. Richard J. Ivanetich of the Institute for Defense Analyses.

Contents

Figures

Tables

Executive Summary

Background

In recognition of the importance of high performance computing (HPC) to maintain and increase the technological superiority of its warfighting and support systems, the Department of Defense (DoD) was directed by Congress to improve its supercomputing capabilities. The High Performance Computing Modernization Program (HPCMP) was established in 1994. Since FY1996, the HPCMP has fielded a world-class HPC infrastructure that has become a critical element for the DoD Science and Technology (S&T) and Test and Evaluation (T&E) communities. The planning and execution of the HPCMP is the responsibility of the HPCMP Office (HPCMPO).

Each year the HPCMP has systematically reviewed its program to assure that it was meeting the requirements of its user community. The Program Review 2000 (PR2000), held during the week beginning on Monday, 27 March 2000, was the first formal and in-depth program review of the HPCMP.

This report presents the author's independent set of observations of the material presented at PR2000. It is not an in-depth review of the program but highlights the important issues—as perceived by the author—that were identified.

HPCMP Organization

The HPCMP comprises five high-level components that provide (1) hardware and operations, (2) intellectual support, and (3) networking and security.

The *hardware and operations components* that provide the computational capability are the Major Shared Resource Centers (MSRCs) and the Distributed Centers (DCs). The MSRCs also administer the subcontracts with universities for the Programming Environment and Training (PET) Program.

Intellectual support comes from the Common HPC Software Support Initiative (CHSSI) and the Challenge Program as well as from the PET Program. The CHSSI supports the development and transitioning of domain-specific software and algorithms into the various HPC systems. The projects selected by the Services in the Challenge Program are their most important projects requiring unusually large amounts of computing.

The Defense Research and Engineering Network (DREN) provides the broadband communications *network,* linking the S&T and T&E sites. It is also responsible for *security.*

These components are complementary and synergistic, all critical to the many successes of the program. Each component is briefly discussed in the following sections.

Major Shared Resource Centers

The four MSRCs established to support HPC for all of DoD have acquired and sustained state-of-the-art equipment in an appropriate operating environment. The facilities, each with a unique organizational structure and functional focus, bring together the talents of large communities of participants, including government agencies, contractor integrators, academic partners, and users.

Findings. Among the accomplishments of the MSRCs are the increases in aggregate computational capacity, mass storage capacity, scientific visualization resources and support, and internal network bandwidth. A new initiative, the Cross-MSRC Collaboration, is underway: its goal is the further sharing of computing capability and data across all the MSRCs—

and ultimately across all HPC computing in DoD.

The facilities at each of the MSRCs are well managed. The effect of the MSRCs has been considerable, primarily in enabling the S&T and T&E communities to support the war-fighter and in establishing DoD as a notable participant in the HPC community.

The MSRCs have matured substantially and are cooperating very much better today than they did a few years ago. However, there continues to be quite independent acquisition planning processes and operational scheduling processes. It is highly desirable to establish an overall strategic planning process across the HPCMP. Similarly, some unification, improvements, and greater flexibility are required for the scheduling processes at the four MSRCs.

Finally, more classified computing resources, currently estimated to be at the 10% level, are needed despite the budget constraints.

Distributed Centers

In addition to the MSRCs, two groups of DCs that focus on more specialized classes of problems have been established: (1) real-time centers that provide support to in-the-loop real-time weapons simulations and experiments; and (2) support centers that provide high-capacity computing facilities supplementing the MSRCs.

Findings. The performance of the DCs as a whole has been good. The support centers provide a good deal of HPC cycles to the user community. The real-time centers are providing a myriad of important HPC services to the community, especially for T&E. They make resources available and enable activities that heretofore were unavailable to the T&E community, and encourage collaborations between the S&T and T&E communities.

Managing the very diverse set of 17 DCs at the HPCMP level is a demanding responsibility and considering the magnitude of the problem, the HPCMPO management team does quite well. The management structure is quite diverse among the DCs, with the result of uneven quality in many cases. Some centers are poorly or under-managed; consequently, they are less effective than they could be.

Good progress has been made in improving the selection process of new DCs and in the oversight process. However, the HPCMPO management is stretched too thin, and the program could benefit from a larger staff able to pay closer attention to each of the DCs. It is important to increase the involvement of the CHSSI, PET, and other programs with the DCs.

There is no procurement funding for DCs in FY2001 and levels of such funding remain uncertain thereafter. Continued efforts to remedy this problem are encouraged.

Networking and Security

DREN connects the MSRCs, the DCs, and remote users with a broadband wide area network. The approach relies on a turnkey commercial services contract.

DREN also is responsible for implementation of security measures for the HPCMP. This aspect of the program is becoming more critical and expensive with time.

Findings. The management of the DREN has been excellent. Here is an example of an element of the HPCMP that is run very successfully directly out of the HPCMPO with the assistance of a small group of very competent contractors, a Technical Advisory Panel, and a number of working groups. One reason for this success is that the technical and operational specifications for the network were established to optimize the network for the program requirements without outside restrictions. This has been both operationally and cost effective.

DREN has also developed collaborative relationships with the Defense Information Systems Agency (DISA), the Ballistic Missile Defense Organization, the Department of Energy, and others to share traffic load or leverage new security approaches—all of which is to the advantage of the HPCMP—and ultimately to our national security.

As the DREN contractual agreement with the commercial supplier is coming to an end, DISA is defining the parameters for the next version of the Defense Information System Network (DISN), which may include provisions for the follow-on DREN. There are serious questions concerning how well DREN, a research and development (R&D) network, might fare if it is included within the specifications of an operational based DISN. The performance, bandwidth availability, risks and costs of the various options, and the relationship with DISA are issues that need very careful consideration. The recognition of DREN by DoD as an independent R&D network, is a possible and reasonable outcome.

Common HPC Software Support Initiative

The goal of CHSSI is to identify, develop, and transition important HPC applications programs to scalable, parallel computing environments. CHSSI effectively leverages HPCMP funding with personnel commitments from the participating Service laboratories/facilities and is an effective mechanism for catalyzing collaborative efforts and work products across DoD.

CHSSI projects are categorized into one of ten Computational Technology Areas (CTAs). The leaders of these CTAs provide economic and requirements analyses, benchmarking, and advice to the HPCMP.

Findings. The management of the CHSSI Program is currently very good. The CHSSI Program Manager attempts to broker and assemble coalitions to improve the effectiveness, the focus, or the range of proposed CHSSI projects, and to improve the evaluation and selection process.

The performance results presented by more than 40 projects spanning the 10 CTAs were quite impressive. The CHSSI programs have accelerated the effective use of scalable parallel systems, contributing to the many successes enjoyed by the S&T community, and to a lesser extent, the T&E community. The most important aspect of this program is its role as facilitator and catalyst, leveraging across DoD and other communities.

CHSSI is a major success despite issues arising from experiential and cultural differences among the Services and especially the S&T and T&E communities, resulting in an imbalance of support—in favor of S&T. Expanding the CTA categories to include more mission-oriented computational areas could help T&E users.

Another issue is how best to "package," distribute, document, and maintain software that has been developed in a CHSSI project. Doing a good job in this technology transfer will reap important benefits.

Programming Environment and Training

The goals of the PET program are to enhance the productivity of the DoD HPC user community; improve the computing and programming environment; transfer innovative HPC technologies into the user community; facilitate collaborative activities with academia, government, and industry; train the user community; and engage partner historically black colleges and universities/minority institutions (HBCU/MI) in joint training and project activities.

Findings. Initial cultural differences between the academic world and the DoD and contractor integrator worlds have been bridged. A cadre of academic, CTA, and computer science experts has been successfully integrated into each of the MSRCs. However, the duality of management—DoD vs. the integrator contractor—continues to be problematical. The PET Program has enhanced the HPC user community through improvements to the programming environment, new and/or improved tools and techniques, and the brokering of collaborative activities with other U.S. government agencies, industry, and academia. This results in enhanced productivity of the user community.

The PET Program has matured nicely and has integrated well into the structure of each

MSRC. Its association with the DCs and much of the T&E community remains limited. All parties would benefit from more PET outreach.

A management decision still to be made is whether or not to unbundle the PET contract from the primary MSRC integrator contracts. The concept of a single PET entity contracted directly to the HPCMPO should be considered. In either case, the PET Program needs a more unified strategic direction, stronger oversight, a greater level of coordination among the MSRCs, and increased involvement with the DCs.

Requirements, Resource Allocation, and Utilization Processes

The *requirements process* is an annual data gathering survey that is useful not only for projecting future resource needs, but also for optimizing the use of all computing assets. The *allocation process*, must reconcile the allocation of the limited HPC resources with the requirements projection. The goal of the *utilization process* is to ensure that DoD's most important HPC needs are satisfied in a prioritized, orderly, and timely manner. Highest priority is given to Challenge projects that are mission critical and high-priority projects with very high computational demands and the capability to efficiently use larger HPC systems.

Findings. The requirements gathering process is serving the HPCMP well. It has an important secondary effect of identifying related projects and brokering coordination and/or cooperation to the benefit of all parties. While somewhat burdensome to the users, it is still worth spending time, money, and effort for this process.

The allocation process is very important—and very difficult. The allocation of HPC resources are, for the most part, fair. But reallocating resources, based on usage demographic and new requirements, is still a clumsy process. Also, it appears that at most Centers, the Challenge projects are using substantially more of the HPC resources than the 25% goals that have been set.

The requirements and facility usage data gathered should be marshalled to argue against calls for "fee-for-service" approaches to fund HPCMP. That will certainly kill the program, thereby frustrating the intent of Congress.

The resource allocation scheme would become more effective with a refinement and standardization of the Service/Agency Approval Authority processes.

Finally, relaxing the high priority of some Challenge jobs might improve the effectiveness of other jobs and make it possible to increase the number of projects that are designated as Challenge projects.

Concluding Remarks

Overall, the HPCMP is well conceived and has developed into a mature, functional, and valuable resource for DoD. Its collaborative efforts and leveraging with the Services, other DoD agencies, U.S. government organizations, industry, and universities have made it a critical element in DoD's ability to insert new technology into weapons systems for increased capability and effectiveness.

And the HPCMP must go on to support the S&T and T&E communities into the 21st century, thereby maintaining this momentum—if given the manpower, money, and freedom to do so. It is critical for the S&T and T&E communities to exert continuing major efforts to connect the relevance of HPC to the many ways in which it supports the war fighter, thereby elucidating the importance of the HPCMP to the Service leadership.

1. Introduction

High Performance Computing (HPC), the high end of the rapidly evolving computer technology, has played and will continue to play a major role in the ability of the United States to maintain and increase the technological superiority of its warfighting support systems. Very high fidelity modeling and simulation are the keys to rapid insertion of advanced technology into defense acquisition, test and evaluation, training and operational activities at every level and in many ways. In 1992, Congress directed the Department of Defense (DoD) to improve its supercomputing capabilities, and in 1994 the High Performance Computing Modernization Program (HPCMP) was established to support DoD Science and Technology (S&T). In 1995, Congress directed that the HPCMP should also support Test and Evaluation (T&E) activities. Since FY1996, the program has fielded a world-class HPC infrastructure that has become a critical element for the DoD S&T and T&E communities. To this day there continues to be strong Congressional interest in the HPCMP. The HPCMP has demonstrated substantive success in making HPC available to a wide-ranging DoD S&T and T&E community.

The mission of the HPCMP, when it was established, is summarized in the program's original formal mission statement.

> "To modernize the total high performance computational capability of the Department of the Defense science and technology and test and evaluation to a level comparable to that available in the foremost civilian and other government agency research and development environments."

1.1 HPCMP Organization

The major organizational components of the HPCMP are illustrated in Figure 1 below. Each of the higher-level components are briefly discussed in the following paragraphs.

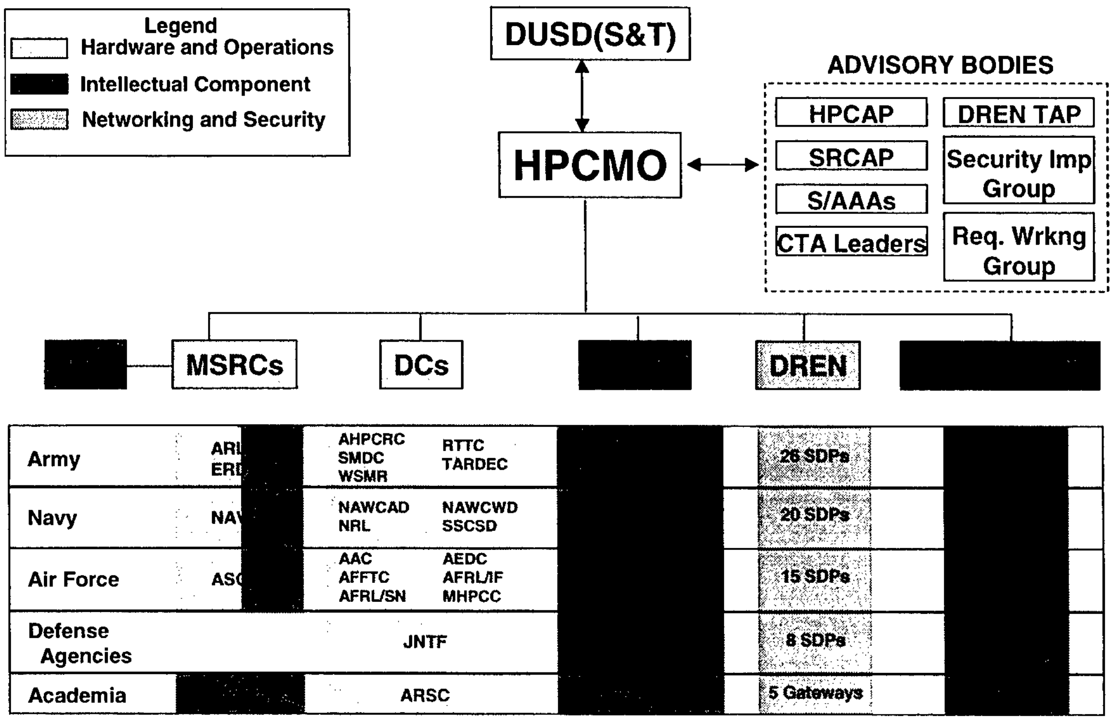

Figure 1. HPCMP Organizational Components

The computational capability that the HPCMP has installed to support the needs of the DoD are distributed amongst the two classes of HPC resource centers:

- Major Shared Resource Centers (MSRCs), and
- Distributed Centers (DCs).

The four MSRCs are funded entirely by the HPCMP through contracts with integrator commercial firms. Only the acquisition of the required hardware is funded by the HPCMP for each of the 17 DCs. The MSRCs also maintain subcontracts with university entities for the Programming Environment and Training (PET) Program. This program, which places academics at each of the MSRCs and maintains connectivity with a larger

pool of computer and domain experts at the universities, serves as an intellectual catalyst for HPCMP.

Two components of the HPCMP contribute heavily to the intellectual fabric of the program and to its ultimate success:

- The first of these is the Common HPC Software Support Initiative (CHSSI). This program supports the development and transitioning of domain specific software and algorithms onto various HPC systems.
- The second is the set of designated "Challenge projects." These are selected by each Service as their most important HPC projects requiring unusually large amounts of computing.

Another important element of the HPCMP is the Defense Research Engineering Network (DREN) which provides the broadband reliable communications network linking the designated sites of the HPCMP S&T and T&E communities. Requirements specifications and operation of this network are under the direct control of the HPCMP, by mutual agreement with the Defense Information Systems Agency (DISA).

These program components are complementary and synergistic. The five structural elements of the HPCMP are all critical to the success of the program. Together they make available a variety of high-end computational resources, with the necessary ancillary and operational support entities and the intellectual capital from both within and outside DoD to take advantage of these resources, in order to support a wide range of S&T and T&E activities. MSRCs, DCs, and DREN make available the necessary hardware and operational resources. The PET component of the MSRCs, CHSSI, and the Challenge projects, support the education/training/tutorial programs, the development and distribution of software, and the effective use of the HPC facilities to solve important but difficult problems for DoD.

The user community is very diverse and includes participants from all the Services, from the S&T and T&E activities, and collaborating participants from academia and industry. In FY99, there were about 5,250 identified users.

The planning and execution of the HPCMP is the responsibility of the HPCMP Office (HPCMPO) which was established in 1994 and reports to the Deputy Under Secretary for Defense for Science and Technology (DUSD(S&T)). In addition to establishing and maintaining the HPC centers, the HPCMPO has defined and refined the functions of the several components that make up the HPCMP. It is assisted in this role by a number of advisory bodies and working groups. The office also continuously surveys the changing (mostly growing) requirements of the community, oversees allocation of the HPCMP resources based upon Service priorities, manages the operations at the major shared resource centers, sets policy for all program components, and plans annual technology insertion acquisitions for the timely evolution of the HPC resources.

1.2 Program Review 2000

Since its inception, the HPCMP has systematically reviewed its program to assure that the program was meeting the requirements of the user community and, more broadly, those of DoD as well. As the HPCMP has matured and its organizational structure and program have evolved, the mechanisms for the review process have changed. Early on, the HPCMP User Group Conferences served as the vehicle for such reviews. With time, the organization has developed more formal focused events for the review process.

The HPCMP Program Review 2000 (PR2000), held during the week beginning on Monday, 27 March 1999, was the first formal, in-depth program review that had been planned for the HPCMP. The primary reason for PR2000 is that it is an appropriate mechanism for periodically reviewing the HPCMP. It brings together in a formal manner the HPCMP management team from across the country along with representative user participants. And because it is an extended activity focused on evaluating all aspects of the program, it

enables the development of an overall critique of the program. In particular, the benefits of such a periodic program review are the gaining of insights into the program with:

- understanding of costs and benefits,
- review and evaluation of program processes,
- identification of issues and lessons learned, and
- insights for targeting future activities.

PR2xxx is planned to be a bi-annual HPCMP event to assess the program's process in meeting its goals and objectives.

Recent laws passed by Congress[1] now require a more formal and structured review process for all programs with information technology (IT) components and/or assets. The PR2000 also serves to satisfy these requirements.

The PR2000 process established a Review Board that primarily was composed of members of the HPC Advisory Panel (HPCAP) or their delegates. The HPCAP is composed of S&T and T&E executives from each of the Services and representatives from selected Defense Agencies. The Review Board was chaired by the Director, HPCMP. Sessions were held along HPCMP organizational lines and were chaired by the HPCMPO Project Managers. Presentations were given by the four MSRC Directors and a sampling of the DC Directors or their delegates. For the resource centers, including the PET Program, these presentations were considered to be part of the ongoing Post-Deployment Evaluation and Assessment Process (P-DEAP). The presentation on Networking and Security was given by the Program Manager. Details of the CHSSI program were presented by the Computational Technology Area (CTA) leaders. The PET Program was described by

[1] Including the Government Performance and Results Act (GPRA) of 1993, the Clinger-Cohen Act of 1996, and DoD regulations such as the ASD(C3I) Guide for Managing IT as an Investment and Measuring Performance.

some of the PET leaders from each of the MSRCs. Finally, the Requirements and Allocation Processes, including the Challenge Projects implementation, were described by the representatives from the Services and Agencies.

Categories to be used for the PR2000 Review for each of the HPCMP components were established. These were

- Management,
- Part performance, and
- Impact on DoD.

The Board members were asked to evaluate each program component with a *pass*, *fail*, or *exceed* rating. A computer based evaluation tool was made available to the Panel members to capture their comments and initiative ratings.

At the conclusion of the presentations, the Project Managers each delivered Summary Out-Briefs. These included a summary of each of the component projects and also their assembly and interpretation of the comments and evaluations made by the members of the Review Board. The HPCMPO maintains a file of all of the briefing material presented and of the commentaries of the Review Board.

This report represents an independent set of observations of the material presented at the PR2000, as requested by the HPCMPO. It is not meant to be an in-depth review of the program, but rather an attempt to highlight important issues that were explored during the course of the week-long review as perceived by the author.

2. HPC Centers—MSRCs

Early in the life of the HPCMP, the concept of establishing and maintaining a number of Major Shared Resource Centers to support high performance computing for all of DoD was developed. During FY96, four MSRCs were established, each with an integrating contractor responsible for the acquisition and operational activities for the center chosen by competitive bid. Table 1 lists the MSRCs and their integrator contractors (both with their current organizational names).

Table 1. Major Shared Resource Centers

Label	MSRC Name	Integrator
ARL	Army Research Laboratory	Raytheon Systems Corp.
ASC	Aeronautical Systems Center	Computer Sciences Corp.
ERDC	US Army Engineer Research and Development Center	Computer Sciences Corp.
NAVO	Naval Oceanographic Office/MSRC Program	Logicon, Inc.

The MSRCs have acquired and sustain a collection of state-of-the-art equipment in an appropriate operating environment. In addition to the basic HPC equipment configured for the class of problems to be solved, they also maintain leading-edge, hierarchical mass storage systems, visualization programs, and internal high performance networks. It should be noted that the HPCMP equipment resources are state of the art but are functional systems chosen to be effective for DoD problems.

Most important, however, the MSRCs bring together the talents of a large community of people—including government, contract integrator, academic partners (PET), and user participants. This crucially important intellectual component includes people with do-

main-specific scientific and engineering skills (computational scientists and engineers), computer expert skills (computer scientists and engineers), and complex facilities management skills. This human resources component, very largely software skill based, is absolutely crucial to the success of any HPC program to support the goals of the functional mission organizations (i.e., the Services).

During the course of the PR2000, successes of the MSRCs in growing HPC computing resources were presented, including:

- aggregate computational capacity,
- mass storage capacity,
- scientific visualization resources and support, and
- internal network bandwidth.

Also, dozens of contributions were credited to the MSRCs for DoD mission success stories.

A major initiative currently underway is a Cross-MSRC Collaboration. Four collaborative projects have been identified, with each of the four MSRCs taking the lead in one of the initiatives. The goals here are to lead in the development of approaches to further the sharing of computing capability and data across all the MSRCs—indeed all HPC computing in the DoD. Indeed these collaborations are facilitating a more general cooperative team-work approach across the MSRCs. The four collaborative projects with the lead MSRC are:

- Security–ARL
- Information Environment–ASC,
- Metacomputing–ERDC
- Mass Storage–NAVO

2.1 Findings

Management. The MSRCs have matured substantially since the establishment of the HPCMP. Internally each MSRC has developed smooth working relationships amongst the various participants including government and integrator contractor personnel, the CTA leaders, the user community, and the academic participants supporting the PET program. Furthermore, a tendency towards competition between the MSRCs with little regard to the state of the total HPCMP has evolved to a cooperative stance, still striving to be individually the best but with the ultimate goal being the success of the HPCMP.

The MSRC directors, along with the HPCMPO management are in constant group communication, by e-mail, a weekly M-bone video-style conference, and with submitted formal quarterly and annual reports. In addition to the usual user-management interactions, the user community also presents its input via the Annual User Satisfaction Survey. The Shared Resource Centers Advisory Panel (SRCAP) could play a more effective role than it is now in facilitating organized feedback from the user community. And the HPCMPO management has a formal semi-annual site visit to each of the MSRCs. One important unifying activity is the set of recently defined Cross-MSRC Collaboration projects. These appear to be very effective in fostering cooperation amongst the MSRCs.

The several acquisition steps defined in the original integrator contracts, allowing for Performance Level (PL) expansions, have now all been executed—PL1 through PL3b. These expansions have taken place with increasing sophistication, with better understanding of the users' requirements, trends in the direction of technology growth, the state of the marketplace, and with the improved use of metrics to evaluate the computing options. Input from all stakeholders has been an important part of the process.

The facilities at all the MSRCs are well managed and information appears to be abundant to keep the user community informed as to the facilities available, the state of user accounts, and most other information of use to a user. Over the life of the MSRCs the inter-

nal infrastructures have grown more robust and the user environment keeps on improving.

Each MSRC has its own scheduling algorithm that appears to complicate the use of the facilities for some users. Most, if not all of them, tend to favor certain job classes over others in a static manner, so some classes of jobs may "never" run. There are better approaches to scheduling. All the MSRCs, however, do appear to be able and willing to modify the job queues by human intervention, if necessary, when it appears sensible to do so.

Performance. Over the several PL expansions, the computation level of performance of the MSRCs have improved substantially, not withstanding the fact that the projected needs tend to exceed the reality of what is possible within the fiscal limits. In parallel with the growth of capacity, there has been good growth in the ancillary facilities, including addressable memory, hierarchical external storage, visualization facilities, and internal and external network bandwidths. These all give rise to an improved capability for each of the MSRCs.

The successes in establishing the DoD HPC Centers as a notable participant in the U.S. and world HPC community are noteworthy. With these accomplishments come important leveraging of the support of the HPCMP for important DoD problems arising from collaborative and sharing efforts with other United States Government (USG), industry, and academic institutions.

Impact. The four MSRCs hosting the major DoD HPC resources serve as the basis upon which the whole HPCMP impacts DoD. The effect has been considerable, primarily in enabling the S&T and T&E communities to support the warfighter. Both communities fully understand the importance of these Centers in facilitating their studies and tests in fulfilling their missions for the Services. Each of the MSRCs has an impressive array of HPC and ancillary resources. With such facilities, the DoD community has fulfilled the original mission of the HPCMP, and made possible relationships and collaborations with

the other advanced HPC communities in government, industry, and academia. This, in turn, increases the contribution of the HPCMP to DoD programs.

Each of the MSRCs has a unique organizational structure and functional focus. Their collective contributions each year are numerous, among them the important contributions directly related to a warfighting need, especially during times of crisis such as in Kosovo and Bosnia. There are many more examples where the direct impact is more of an enabling of the S&T and T&E communities, with the impact on the warfighter being more indirect or with a longer time constant but of no less importance.

2.2 Discussion

The MSRCs have matured very nicely since the start of the HPCMP. Their acquisitions have been very effective and have become more so with time. However, there is concern that the acquisitions have not been coordinated across the whole program as well as they might have been. *The establishment of an overall strategic planning process across the HPCMP is highly desirable.* Here the HPCMPO might facilitate the difficult tradeoff of architectural diversity vs. most cost-effective computational cycles. Having each MSRC specialize in a particular architecture – spanning all viable architectures would be highly desirable.

Operationally, although all the MSRCs appear to document their facilities well, it does appear that some members of the user community seem not to be able to access the facilities as effectively as they expect. One example of such a difficulty is the different job queue handling processes and scheduling policies across the MSRCs and across platforms even within a single MSRC.

Another problem is the limited availability of classified computing resources, currently estimated to be at the 10% level. Currently, only ARL has substantial capability to support such computing. There clearly is pressure for more assets to be made available for

classified computing. With the current stringent budget situation, it is difficult to relieve this problem, but it represents a real requirement that should be addressed.

Continued effort is encouraged to foster collaboration amongst the MSRCs. Much progress has been made but there is much more that can be done. The eventual goal ought to be an approach to a single DoD metacomputing entity as perceived by a user.

3. HPC Centers—DCs

From the beginning of the HPCMP, it became clear that there was a need for a set of HPC resource centers more widely distributed than the MSRCs that would focus on more specialized classes of problems. Two groups of Distributed Centers (DCs) have been established by the HPCMP. One group is composed of primarily "real-time" or near "real-time" centers providing the capability to support in the loop real-time weapons simulations and experiments. Many of these are located close to or at T&E centers, and are particularly valuable to the T&E community; they also serve the S&T community. A second group serves as support centers providing high capacity computing facilities complementing and supplementing the MSRCs.

Unlike the MSRCs, operational activities at the DCs are not normally supported with funds from the HPCMP. The HPCMP does support the acquisition of the required HPC hardware and basic software assets for the DCs. And the DCs tend to be substantially smaller in total computational capability than the MSRCs. With this arrangement, the Services effectively share costs with the HPCMP in the support of the DCs. It should also be noted that some of the DCs have been established and continue to be supported by ongoing Congressional interest. In some of these cases, Congress has directed Research, Development, Test and Engineering (RDT&E) funds to be used to support operations of these centers.

Currently, 17 DCs have been established; these are listed in Table 2. Over the last few years, including FY2000, the procurement budget for DCs has been about $19 million per year. For FY2001, the guidance is for no funds for enhancement to the existing DCs or for new DCs. And POM02 guidance for future year DC procurements is also disturbingly low.

Table 2. Distributed Centers

Label	DC Name	Location
AAC	Air Force Air Armament Center	Eglin AFB, FL
AFFTC	Air Force Flight Test Center	Edwards AFB, CA
AFRL-IF	Air Force Research Laboratory/Information	Rome, NY
AFRL-SN	Air Force Research Laboratory/Sensors	Wright-Patterson AFB, OH
ARSC	Arctic Region Supercomputing Center	Fairbanks, AK
AHPCRC	Army High Performance Computing Research Center	Minneapolis, MN
AEDC	Arnold Engineering Development Center	Arnold AFB, TN
JNTF	Joint National Test Facility	Schriever AFB, CO
MHPCC	Maui High Performance Computing Center	Kihei, HI
NAWC-AD	Naval Air Warfare Center Aircraft Division	Patuxent River, MD
NAWC-WD	Naval Air Warfare Center Weapons Division	China Lake, CA
NRL-DC	Naval Research Laboratory	Washington, DC
RTTC	Redstone Technical Test Center	Redstone Arsenal, AL
SMDC	Space & Missile Defense Command	Huntsville, AL
SSCSD	Space and Naval Warfare Systems Command Center	San Diego, CA
TARDEC	Tank-Automotive Research, Development and Engineering Center	Warren, MI
WSMR	White Sands Missile Range	White Sands, NM

Because an in-depth review of the 17 DCs is not possible nor desirable to perform every year, the HPCMPO has chosen to review only 5 of the DCs for PR2000. The five DCs reviewed in PR2000 were those that were funded in FY98:

- Air Force Research Laboratory—Rome, NY (AFRL/IF)
- Naval Air Warfare Center Aircraft Division—Patuxent River, MD (NAWC-AD)
- Space & Missile Defense Command—Huntsville, AL (SMDC)
- Space and Naval Warfare Systems Command Center—San Diego, CA (SSCSD)
- White Sands Missile Range—White Sands, NM (WSMR)

In subsequent year PR2xxx reviews, other DCs will rotate into review status.

3.1 Findings

Management. By their nature, the DCs represent a very diverse set of HPC resources satisfying a wide range of needs. Managing 17 DCs at the HPCMP level is a demanding responsibility. The management difficulty both at the HPCMPO level and at the individual DC level is further complicated by the fact that no funds are made available by the HPCMP for operations at these centers (with the exception of those DCs where Congressionally mandated funds have been made available). And there are a number of examples of DCs that have not achieved their acquisition goals.

At the HPCMPO level, there is in place an effective process (mostly annual) for the selection of new DCs. This process includes a call for proposals with well-defined criteria, a formal validation and review process, a technical review, and finally a setting of priorities against DoD needs and a match against fiscal constraints. The process has been improved over the years and now includes site briefings, review of past performance of the proposing entity, proposal performance metrics, etc. At this level, oversight of the ongoing operations and facility utilization occurs with periodic reports, daily e-mail, and telephone interactions, biennial site visits, P-DEAP assessments of selected DCs, and

phone interactions, biennial site visits, P-DEAP assessments of selected DCs, and independently by direct customer feedback. Considering the magnitude of the management problem, the HPCMPO management team does quite well. However, the program could benefit from a larger staff able to pay closer attention to each of the DCs.

At the individual DCs, the management structure appears to be as diverse as are the missions of these centers. Each has established an infrastructure that fits into the constraints of its location and best meets the needs of its missions. Those DCs, primarily the support centers, with HPCMP funding for operations, appear to be well managed. For the real-time centers, there appears to be a wide range in the capabilities of management. In many cases, these centers are under-managed or managed by inexperienced managers; consequently, some are less effective than they might be otherwise.

Performance. The performance of the DCs as a whole has been extremely good. The support centers provide a good deal of capacity HPC cycles to the user community in a very effective manner. The real-time centers are providing a myriad of important HPC services to the community especially for T&E. They make available resources and enable activities that heretofore were unavailable to the T&E community. Furthermore, their existence serves as a natural path for members of the T&E community to enter the HPC realm, become familiar with what is available, and to take advantage of the physical and intellectual assets within the broader community. This encourages collaborations between the S&T and T&E communities with both communities benefiting from the relationship.

Impact. The effect of the DCs has been as diverse as are the DCs themselves. These centers have supported a very wide range of both S&T and T&E functions, with direct influence on activities that support the warfighter. More generally, this facet of the HPCMP has introduced a very broad set of participants into the HPC culture, with a resulting synergy and leveraging of the physical and intellectual resources across DoD and including other U.S. Government agencies, industry, and academia.

3.2 Discussion

The DCs have contributed in important ways to the success of the HPCMP. And good progress has been made as the program has matured with respect to the selection process, the diversity of the various DCs supported, and the oversight process established to encourage high levels of performance. The management team in the HPCMPO is excellent—but this facet of the HPCMP is very demanding and the management team is probably stretched too thinly. Closer involvement of the HPCMPO, more frequent visits, and a greater involvement amongst the various players would likely pay handsome dividends.

There is very little involvement by the PET Program in the DCs, leaving an intellectual vacuum. There is an opportunity here to better leverage the resources if one could generate a greater level of interaction and/or collaboration with other DCs and the MSRCs, especially through the PET and CHSSI Programs. Developing relationships across these entities will certainly enhance the contributions made by the DCs and ultimately their impact on the warfighter.

The lack of procurement funding in FY2001 and uncertain levels of such funding thereafter are certainly disturbing. Continued efforts to remedy this problem are encouraged. Also, the issues of continued replenishment of resources for the DCs and/or controlled "sunsetting" of some are financial-based issues that require some study and planning.

Another difficult issue is how to encourage better management at the DCs. Without identified operational funds, this leaves a managerial void that frequently leads to inexperienced personnel filling the gap.

4. Networking and Security

The Defense Research and Engineering Network (DREN) first came on line in mid-1997. It now connects the MSRCs, the DCs, and remote users with wide area network (WAN) data transport at appropriate bandwidths. These bandwidths range from fractional T3 (45 million bits/second) to OC3 (155 million bits/second) with near-term future capabilities rising to OC12 (620 million bits/second) and beyond. DREN continues to grow in the number and type of nodes and aggregate bandwidth capability to satisfy the growing needs of the S&T and T&E communities. Today there are over 75 sites connected.

Because this network supports S&T and T&E research and engineering activities and not command and control or other operational activities, its technical specifications and requirements are different than those required for most DoD networks. Therefore, with DISA concurrence, the HPCMPO deals directly with the network contractor on all technical issues. Thus the HPCMPO has the capability to monitor DREN activity and to oversee changes to meet the mission requirements.

In addition to networking support, DREN also is responsible for implementation of security measures for the HPCMP. This aspect of the program is becoming more critical with time and will continue to do so for the foreseeable future. Simultaneously, the cost of supporting this aspect of the program continues to rise. DREN also maintains working relationships with a number of other U.S. government organizations involved in testing new extensions in networking and security.

4.1 Findings

Management. The management of the DREN, security, and network security for the HPCMP has been superb. Here is an example of an element of the HPCMP that is run

very successfully directly out of the HPCMPO with the assistance of a small group of very competent contractors. The management is supported by a Technical Advisory Panel and a number of working groups all populated with appropriate experts from the Services. The approach works very well and the results are highly appreciated by the HPCMP community.

The approach has been to rely upon a turnkey commercial services contract that allows for bandwidth to be installed anywhere in the country, provides high scalability, and allows for the insertion of new technology—all on very short notice. Important management chores include assessing the requirements for new sites and requirements for bandwidth upgrades. Because the demand for resources keeps growing, the DREN management team must make the difficult decisions on how best to support all the HPCMP needs, or alternately, who not to serve.

On the security front, the approach is for "defense in depth." A wide range of technical and user-access policy approaches has been taken to ensure confidentiality, availability, accountability, and the integrity of user information have been instantiated. Network intrusion detection systems are in place.

A number of collaborative relationships have been developed with DISA, Ballistic Missile Defense Organization (BMDO), the Services, Department of Energy (DOE), and others. Examples include the sharing of traffic load and the development of new security approaches. These teaming/partnership efforts give use to a leveraging of resources to the advantage of the HPCMP.

Performance. The networking and security program has made available the required networking and security resources to the HPCMP in a timely manner. Performance has grown from a T3 (45Mbps) backbone Interim Defense Research and Engineering Network (IDREN) in FY1997 to an OC3 (155Mbps) backbone, growing to OC12 (620 Mbps) in today's DREN. And the extent of the network has grown to accommodate more and more of the distributed user base.

Similarly, as the demand grows for more secure procedures within a node and between nodes, the HPCMP has risen to the occasion in a very natural way.

Impact. The impact of DREN on DoD has been enormous. It has demonstrated what a high performance network is capable of accomplishing. The MSRCs, the DCs, numerous geographically dispersed DoD laboratories and test centers, as well as industrial and academic sites, are interconnected with high bandwidth reliable networked interconnections. This makes available specialized assets, especially HPC, to a wide community whose mission ultimately is to support the warfighter. It facilitates new important concepts such as meta-computing and access to off-site mass storage—concepts that are in the process of becoming very important to DoD.

In other dimensions, it has enabled members of the HPCMP to participate in leading-edge technology demonstrations and test-beds. These, in turn, lead to effective and early transfer of new technology into operational networks. Similarly, in the security arena, the HPCMP has served as a leader placing in test some of the new methods in information technology security practices.

4.2 Discussion

DREN has been very effective; its management is first rate. DISA, occupied with more urgent command and control issues, chose not to involve itself in the acquisition of the DREN. This gave HPCMP managers opportunities to optimize the network for the program requirements without outside restrictions. This has been both operationally very effective as well as being quite cost effective. The DREN has established itself as one of the most effective components of the HPCMP. DREN meets the leading-edge communications requirements of the DoD HPC R&D community. As the HPCMP further matures, DREN will be expected to provide yet more bandwidth to the user community. Additional funds will likely be required to satisfy these needs.

The contractual agreement with the commercial supplier is coming to an end now and a new contract must be let or an option to extend the current contract for a limited time must be executed. DISA is now defining the parameters for the next version of the Defense Information System Network (DISN), which may include provisions for the follow-on DREN. There are serious questions concerning how well DREN, a research and development (R&D) network, might fare if included within the specifications of an operational based DISN. Of particular concern to HPCMP is the management and operational control of DREN and the issues of quality of service (QoS), prioritization, and preemption. Although there is some merit in an economy of scale, it is not at all clear that such disparate needs can both be serviced effectively with one policy set. The performance, risks and costs of the various options and the relationship with DISA are issues that need very careful consideration. The establishment of DREN, recognized by the DoD as an independent R&D network, is a possible and reasonable outcome.

Another issue of interest is that some of the T&E needs are more focused on large data transfers with time limitations rather than on HPC-based requirements. The issue of whether the limited bandwidth resources of DREN should be allocated to T&E needs based solely on bandwidth requirements needs some study. And the projected budget for future year bandwidth acquisition increases only modestly. Whether it will be adequate to meet the rapidly growing needs of the community is unclear—much depends upon the marketplace cost of bandwidth in the near future. If the unit price for bandwidth does not decrease substantially, there will be severe bandwidth limitations for DREN.

5. Common High Performance Computing Software Support Initiative

The goals of the Common HPC Software Support Initiative (CHSSI) is to identify and develop and then to transition important DoD HPC applications programs to a high performance scalable, parallel computing environment. With this initiative, support is provided for projects for the conversion of or the development of domain-specific software, which is efficient, scalable and portable, and/or algorithms, tools, models, and simulations, which have been selected from peer reviewed proposals. Some of the projects currently supported include computer codes for computational chemistry (GAMESS), fluid dynamics (Cobalt 60), and shock physics (CTH).

CHSSI effectively leverages HPCMP funding with personnel commitments from the participating Service laboratories/facilities and is an effective mechanism for catalyzing collaborative efforts and work products across the DoD community. This results in a growing portfolio collection of important HPC applications that can be executed on a variety of HPC platforms and by groups not necessarily those that were involved in the initiative. CHSSI also has the effect of building HPC experience for defense applications among DoD, industry, and academic scientists and engineers, and of leveraging other HPC efforts across DoD, U.S. government organizations, industry, and academia. For the PR2000, the status of those CHSSI projects selected during the period FY1995 – FY1997 was presented for review.

The CHSSI projects are categorized into one of the ten Computational Technology Areas (CTAs) that have been identified by the HPCMPO. These CTAs, along with the current CTA leaders, are listed in Table 3. The CTA leaders constitute the CTA Advisory Panel (CTAAP) for the HPCMP, in which they play critical roles in economic analyses, re-

quirements analyses, benchmarking, and advising. The CTA leaders represent the user community in each CTA to the HPCMPO and at national conferences. They also keep their communities aware of events and issues of relevance to them. They work with other CTA leaders and with the MSRC staffs, the HPCMPO, DoD and Federal Agency HPC managers, and academic HPC experts under the PET effort. This is to ensure collaboration where appropriate, reduce redundancies, and to ensure users gain access to the HPC resources they need. They attend or direct various HPC and CTA related meetings, conferences, workshops, and seminars throughout the year. They also develop the DoD HPC user community through CTA Web page development, success story publications, technical papers, presentations, peer reviews, and refereed journal articles.

Table 3. Computational Technology Areas

Label	Computational Technology Area	CTA Leader/Organization
CSM	Computational Structural Mechanics	Raju Namburu/ARL
CFD	Computational Fluid Dynamics	Jay Boris/NRL
CCM	Computational Chemistry and Materials Science	Leslie Perkins/AFRL
CEA	Computational Electromagnetics and Acoustics	Bob Peterkin/AFRL
CWO	Climate/Weather/Ocean Modeling and Simulation	George Heburn/NRL
SIP	Signal/Image Processing	Richard Linderman/AFRL
FMS/C4I	Forces Modeling and Simulation	Bob Wasilausky/SPAWAR
EQM	Environmental Quality Modeling and Simulation	Jeff Holland/ERDC
CEN	Computational Electronics and Nanoelectronics	Barry Perlman/CECOM
IMT	Integrated Modeling and Test Environments	Andrew Mark/ARL

The CHSSI program currently is funded at about $20 million per year and supports about 45 projects. The project selection process is peer reviewed within each CTA and aligned with DoD priorities. The CTA leaders provide much of the necessary review and oversight effort, and are the action agents for the CHSSI projects. There are now 76 computational codes and library sub-codes that have been parallelized. The Challenge Projects and CHSSI leveraging have allowed application codes to be effectively executed on systems with over 1,000 processors and also to conduct physics-based experiments not possible before CHSSI.

5.1 Findings

Management. The management of the CHSSI Program is currently very good. Serious efforts are in progress to better focus the proposals that are submitted and to improve the evaluation and selection process. The CHSSI Program Manager attempts to broker and assemble coalitions to improve the effectiveness, the focus, and/or the range of proposed CHSSI projects. Attempts are being made to select those proposals that are best aligned with DoD priorities. Special emphasis is now being placed on improving the level of T&E participation in CHSSI. Once projects are chosen, peer and milestone reviews ensure an effective execution of the project and also a good return on the investment. The CHSSI program has been very successful in forging symbiotic teaming partnerships across Service boundaries and leveraging their effectiveness by including other U.S. government, academic, and industrial participants.

Each of the several CTAs work quite differently, consistent with the needs of their natural constituency, closer to basic science, or to weapons systems, or to test and evaluation. These differences affect the partnering arrangements and the focus of each CTA in a quite natural manner. It may be useful to encourage an exchange of lessons learned across the CTAs.

Performance. The performance results presented by more than 40 projects spanning across the 10 CTAs was generally quite impressive. The reporting mechanisms for scal-

ability were not uniform across projects, platforms, number of processors, etc. As is to be expected, the variance in scalability is a function of the class and size of problem, and of the architecture of the target system. In some cases, unrealistically high values for scalability were presented, indicating an inadequate understanding of how to make the measurement. Several CTA areas presented meaningful and good-to-excellent scalability figures. CFD, in particular, presented an understanding that CHSSI performance was more than just achieving a high value for scalability, but also includes good understanding of the computational issues, the technology and the sociology.

Impact. The existence of the CHSSI program has facilitated and accelerated the effective utilization of scalable parallel systems. Consequently, the various CHSSI projects have contributed many important results supporting the S&T and, to a lesser extent, the T&E entities participating in new and/or improved weapons design, acquisition, testing and evaluation over a wide range of activities. CHSSI projects have also contributed to and/or improved computational operational studies including such diverse arenas as ocean and wave modeling, real-time simulation of training exercises, and intelligence applications. However, the most important impact of CHSSI has been to serve as a catalyst for bringing together participants from across the Services and Agencies, from other U.S. government entities, and from industry and academia in both developing and in sharing the resulting applications program. This gives rise to a real leveraging of intellectual resources, across a much larger community than would be possible with a single Service approach, in applying these resources to support the warfighter on many different fronts. One good example of this is that 15 CHSSI-developed codes were used by Challenge Projects to facilitate solutions of those important Service-designated problems.

5.2 Discussion

The CHSSI program has been maturing and improving as the program manager and the participants gain experience. CHSSI must be considered a major success. The current management is effective and is working hard at understanding how to further improve the

initiative. One of the most important priorities is how best to allocate projects so as to be most effective and to focus on the most critical applications.

A number of issues arise, partly from experiential and cultural differences, resulting in the imbalance of support amongst the Services and especially between the S&T and the T&E communities. The S&T community appears to utilize the largest percentage of the resources. This comes about in a natural manner, having to do with the greater experience with HPC and a greater historical reliance on using "sharable" software packages by the S&T community, and also the make-up of the CTAs. The CTAs are defined primarily along S&T intellectual organizational lines. There is interest on the part of the T&E community and there may be some merit in expanding the CTA span to include some more mission-oriented computational areas.

Another issue that requires additional resources or effort is how best to "package," distribute, document, and maintain software that has been developed in a CHSSI project. This technology transfer problem has been an endemic one with software developed by a research community for a very long time—there is no simple solution. But doing a good job here will reap important benefits.

6. Programming Environment and Training Program

The Programming Environment and Training (PET) Program was established to enhance the productivity of the DoD HPC user community. Its goals are to:

- improve the HPCMP computing and programming environment;
- transfer innovative HPC technologies into the user community;
- facilitate collaborative activities with academia, government, and industry;
- train the user community; and
- engage historically black colleges and universities/minority (HBCU/MI) partner institutions in joint training and project activities.

The PET Program was integrated into the original integrator contracts established by each of the four MSRCs during FY96. The funding level was set at and remains at $4.5 million per year. The approach was to place world-class academic teams at each of the MSRCs to facilitate the goals of the program.

As the PET Program has evolved, the four MSRCs have differing arrangements in which their PET Programs work. ASC, ARL, and ERDC have similar models with shared academic team activities organized by the designated local CTA leaders. Typically, there is an onsite academic at the MSRCs, and connectivity with a larger pool of computer and domain experts is maintained at the cooperating universities. At NAVO a "tiger team" approach has been developed. These teams are project oriented and flexible in their makeup and approach to the problems. The universities participating in the PET Program and the MSRCs with which they play some role are shown in Table 4.

Table 4. PET Academic Partners

University	ARL	ASC	ERDC	NAVO
NCSA/University of Illinois	XX	X	X	
OSC/Ohio State University	X	XX	X	
Mississippi State University	X	X	XX	X
Syracuse University	X	X	X	X
Rice University	X	X	X	
SDSC/University of California, San Diego				X
Oregon State University				X
University of Texas			X	
University of Tennessee	X	X	X	
Jackson State University*	X	X	X	
Clark Atlanta University*	X	X	X	
University of Virginia				XX
University of Southern Mississippi				X
Morgan State University*	X			X
Grambling State University*				X
North Carolina A&T University				X
Tennessee State University*				X
Alcorn State University*				X
Central State University*		X		
University of Southern California			X	

* Indicates HBCU/MI Institution

X Indicates University involvement at MSRC

XX Indicates MSRC's Lead Academic Institution

6.1 Findings

Management. To minimize the number of contracts and to limit the level of integration required to initiate the HPCMP, it was decided to bundle the PET program as a subcontract into the original primary MSRC integrator contracts. This did lead to some difficulties early in the HPCMP because the academic culture of the PET subcontractors was quite unfamiliar to the commercial integrator firms. And on their part, the academic community had no understanding of the DoD culture. Consequently, the HPCMP had difficulties early on in not only orchestrating a coherent program among the four MSRCs but also in melding in the PET Program. However, after these early trials, the cultural differences have been bridged, and a cadre of academic CTA and computer science experts has been assembled and successfully integrated into each of the MSRCs. However, the duality of management, the U.S. government vs. the integrator contractor, continues to cause problems. The merits of the program are evident to all participants in the HPCMP. Indeed, PET is a very important part of the intellectual fabric that is critical to the success of the HPCMP.

Performance. As planned, the PET Program extends the operational level of user assistance made available by the MSRC operational staff. PET provides formal and individual training as required, but most importantly it assists in establishing a modern HPC environment. These activities tend to be MSRC centric, with each MSRC focusing on the problem of the local subset of CTAs and the local user community. As the program matures and the technology makes it easier, distance learning programs are increasingly used. Also, as the MSRCs themselves deepen their level of cooperation, so too do the PET Programs.

Impact. The impact of the PET Program has been considerable. In each of the MSRCs, the PET Program has influenced in important ways the programming environment, introduced tools and techniques from other user communities, and brokered collaborative activities with other government agencies, industry, and academia. All of these have enhanced the productivity of the HPCMP user community—and consequently their sup-

port to the warfighter. A good number of CHSSI and Challenge projects also have received leveraging assistance from the PET Program.

6.2 Discussion

The PET Program has matured very nicely and is quite well integrated into the structure at each of the MSRCs. There is only limited association with the DCs and consequently to many in the T&E community who are more likely naturally to participate with a DC rather than an MSRC. Although there is some outreach to the T&E community, it is not the norm nor is it extensive. It should be noted that ARL seems to have been more successful than the other MSRCs in such outreach—partly due to the "lucky" situation that it is located at the Aberdeen Proving Ground which is an Army T&E site. And because the T&E community is relatively new to the HPC culture, outreach to this community could reap important benefits. It is important to encourage such outreach.

As the integration contracts for all the MSRCs are coming to closure points and soon must be extended or renegotiated, the question arises as to whether or not there is merit in unbundling the PET contract from the primary MSRC integrator contracts. Although it might be destabilizing in the first instance, the concept of a single PET contract coupled directly to the HPCMPO is worthy of study. However, it is important to keep the PET participants close to the user, especially the new user community and also to the resource centers. Their role should still be driven by the stakeholders' HPC needs. Whether or not the contractual arrangements are changed, there is merit in the HPCMPO developing a unified strategic direction, exercising stronger oversight, encouraging a greater level of coordination amongst the MSRCs, and bringing PET efforts to the DC user communities. The current HPCMPO management is well suited to take on this difficult task.

7. Requirements, Resource Allocation, and Utilization Processes

In its first year, the HPCMP initiated a HPC requirements study to define the nature and the magnitude of the upcoming resource acquisitions. Starting in 1996, when the requirements gathering was instituted as an ongoing annual update process, the requirements process has been continuously refined. The requirements process today is a bottom-up detailed survey gathering data that is useful not only for projecting future resource needs but also in the allocation process to optimize the use of all the HPCMP computing assets. Only approved and funded S&T and T&E projects are included; and in all cases, the requirements are reviewed and validated by S&T and T&E executives.

The allocation of the HPC resources must be reconciled with the requirements projections, which typically exceed the availability of resources. Furthermore, since FY97, when Challenge Projects were introduced, a percentage (20%–30%) of the HPC resources is dedicated to these important projects (see next paragraph). The remaining resources are allocated on a fair-share basis to the Services and other Agencies. The actual resource allocation by project is determined by designated authorities, the Service/Agency Approval Authorities (S/AAA). The goal is to ensure that DoD's most important HPC needs are satisfied in a prioritized, orderly, and timely manner.

DoD Challenge Projects were instituted on a trial basis in FY97 and on a permanent basis in FY98. These are mission-critical high priority projects with very high computational needs capable of efficiently using the larger HPC systems. A minimum threshold requirement of 25 Gflop-years/year is imposed on these Challenge Projects. They receive the highest possible priorities while they are executing to minimize their turnaround time. In many cases, they require resources across multiple systems and centers. Proposals for

Challenge Projects are nominated and certified by Service executives, and these nominated proposals are evaluated and the Challenge Projects selected by a Defense Challenge Projects Allocation Board.

To maintain order and fairness to the user community and to optimize the utilization of the HPCMP resources, all use of the programs resources is monitored. This information, which is tracked by project, is made available to the user community and is also used by MSRC, DC, and HPCMPO management. The information tracked includes not only cycles used but also time in queue and total turnaround time.

7.1 Findings

Management. Since its initiation, the requirements gathering process has been continuously refined, and the data collected has served very well to plan for new acquisitions and for defining HPC resource allocations. The process is quite intensive and includes detailed questionnaires, face-to-face interviews, and Service executive validation. The process is somewhat burdensome to the users but it is widely believed that it is well worth the intrusion.

Given that the resource demands exceed the available resources, the allocation process is very important and quite difficult. The designation of Challenge Projects identifies the allocation of the first (highest priority), around 25% of the available resources in a very natural manner, thus satisfying DoD's highest priority needs. The allocation of most of the remainder of the HPC resources does a good job of attaining effective total utilization of the resources. However, the adjustment processes in place to reallocate resources based upon actual usage demographics and newly developing requirements are rather clumsy and are in need of improvement. Addressing these issues, some recent changes in the allocation/use policy are very positive. These include the ability to transfer allocations across platforms, the open access to "background" cycles, and the designation of some "reserve" (non-allocated) cycles.

The continued refining of the gathering of usage statistics is to be commended. The more complete and uniform this assembled data is across the whole HPCMP, the more effectively the program will be able to acquire and manage resources. Additional effort is required to make such data available in an easy meaningful, and timely manner in a uniform and easily accessible format to all stakeholders. Providing a HPCMP-wide web-based approach makes a lot of sense.

Performance. The requirements gathering process has been quite effective. In addition to gathering the data, the process has an important secondary effect of identifying related projects and brokering coordination and/or cooperation to the benefit of all parties.

The introduction of the Challenge Projects certainly has improved the computational performance level for those projects. This certainly improves the performance effectiveness of these highest priority projects as identified by the Services.

Impact. The effective understanding of the requirements and the appropriate allocation of the HPCMP resources results in optimally satisfying the programmatic HPC needs for maximum impact on the DoD. However, it appears that at most Centers, the Challenge Projects are actually using substantially more of the HPC resources than the 25% goals that have been set. Some scheduling or job priority adjustments may be appropriate.

7.2 Discussion

An important result of gathering detailed requirements and facility usage statistics is the ability to satisfy many of the oversight requirements demanded by Congressional action. Most importantly, it can be used to argue against calls for "fee-for-service" approaches to force what appears to be cost-effective behavior by fiscal executives. *Fee-for-service is not a viable option when dealing with an expensive resource, where the return on investment may not be seen in traditional "business" terms until years later.* And such an approach will certainly kill the program, thereby actually frustrating the intent of Congress.

Although the resource allocation scheme currently in place works quite well, it would be valuable to refine the S/AAA process. There is great variability in the processes used and in the knowledge and energy levels of the several Service/Agency S/AAAs. Some recent changes, such as allowing resource allocations to be used on any system and the open use of any approved project to compete for background time, are commendable. However, the absolute high priority of all Challenge Project jobs may bear some rethinking. Downgrading that priority for those Challenge Project jobs that are determined by the user to not require a high level of priority might improve the effectiveness of other jobs. With such an approach it may be possible to increase the number of projects that are designated as Challenge Projects. And although there is now a process in place to allocate time to new users or users who have no remaining resource allocation, the current process certainly can be improved.

8. Concluding Remarks

Over the past five years, the HPCMP has responded to the needs of its user community, refined its program, matured its organization and process, enabled real and virtual collaboration within DoD and with other HPC communities, and demonstrated joint cooperation amongst the Services. The user community is growing in breadth, numbers, and maturity. It is learning how to take advantage of HPC facilities to advance their missions. In this respect the S&T community is quite mature and the T&E community is learning quickly; but more support is needed from the HPCMP to facilitate this growth.

The HPCMP is well conceived and has developed into a mature, functional, and valuable resource for the DoD. The HPCMP coherently evaluates the requirements and acquires the specialized HPC resources on a DoD-wide basis to most cost effectively leverage the HPC investment. The four MSRCs acquire their resources spanning the diverse needs of the S&T and T&E communities with the buying power leveraging arising from a coordinated acquisition. The DC acquisitions are focused on the more specialized needs of the community. This is a very good strategy—the corporate investment is paying off.

Most of the Shared Resource Centers' assets are organized in unclassified configurations. There is a need for additional resources to support classified activities. Also, the HPCMP does not have a process in place to evaluate new architectures and new approaches to HPC appropriate for DoD. It is desirable to establish such a process as part of the HPCMP.

The intellectual components of HPCMP reside in CHSSI, the Challenge Projects, and the PET program. Together they make for a very robust mission-oriented program. The software-based activities of the HPCMP are primarily centered in these components. The functional performance and effectiveness gains to be achieved with software improve-

ments are at least as important—probably even more important per unit investment—than hardware investments. Consequently, it may be prudent to better focus the activities and increase support for these components.

The HPCMPO performs jointly to satisfy very well the priorities of the Services. It achieves synergistic leveraging of HPC capabilities by considering both S&T and T&E requirements. It has refined the HPCMP and matured its processes. It should be encouraged to continue its good work. Considering S&T and T&E requirements across all the Services together allows for both tactical and strategic decisions on the allocation of resources to be made more effectively. This could be done by facilitating short-term allocation changes within the entire DoD HPC enterprise to deal with contingencies and opportunities.

HPCMP supports the S&T community well and is working to better support the T&E community. For success, continued effort is required by both the HPCMP and the T&E participants. Here the HPCMP has the opportunity to catalyze symbiotic teaming of the S&T and T&E communities. Both the S&T and T&E communities require continuation of a robust HPCMP program.

There are many demonstrated successes of the HPCMP both in the long and short term. It is imperative that the HPCMP continues to be supported as a world-class HPC resource for the DoD. It is an essential element in the successful DoD adoption and insertion of new technological advances to provide for improvements and increases in specific capabilities. It is most important for the S&T and T&E community to exert continuing major efforts to connect the relevance of HPC to the many ways in which it supports the war fighter, thereby elucidating the importance of the HPCMP to the Service leadership.

Acronyms and Abbreviations

AAC	Air Armament Center [DC]
AEDC	Arnold Engineering Development Center [DC]
AFFTC	Air Force Flight Test Center [DC]
AFRL/IF	Air Force Research Laboratory/Information Directorate [DC]
AFRL/SN	Air Force Research Laboratory/Sensors Directorate [DC]
AHPCRC	Army High Performance Computing Research Center [DC]
ARL	Army Research Laboratory [MSRC]
ARSC	Arctic Region Supercomputing Center [DC]
ASC	Aeronautical Systems Center [MSRC]
ASD(C3I)	Assistant Secretary of Defense for Command, Control, Communications, and Intelligence
BMDO	Ballistic Missile Defense Organization
CCM	computational chemistry and materials science [CTA]
CEA	computational electromagnetics and acoustics [CTA]
CECOM	Communications Electronics Command [U.S. Army]
CEN	computational electronics and nanoelectronics [CTA]
CFD	computational fluid dynamics [CTA]
CHSSI	Common High Performance Computing (HPC) Software Support Initiative

CSM	computational structural mechanics [CTA]
CTA	Computational Technology Area
CTAAP	Computational Technology Area Advisory Panel
CWO	climate/weather/ocean modeling and simulation [CTA]
DC	Distributed Center
DISA	Defense Information Systems Agency
DISN	Defense Information Systems Network
DoD	Department of Defense
DoE	Department of Energy
DREN	Defense Research and Engineering Network
DUSD(S&T)	Deputy Under Secretary for Defense for Science and Technology
EQM	environmental quality modeling and simulation [CTA]
ERDC	Engineer Research and Development Center [MSRC]
FMS/C4I	forces modeling and simulation/command, control, communications, computers, and intelligence [CTA]
FY	fiscal year
GPRA	Government Performance and Results Act
HBCU/MI	Historically Black Colleges and Universities/Minority Institutions
HPC	High Performance Computing
HPCAP	High Performance Computing (HPC) Advisory Panel
HPCMP	High Performance Computing Modernization Program
HPCMPO	High Performance Computing Modernization Program Office

IMT	integrated modeling and test environments [CTA]
IT	information technology
JNTF	Joint National Test Facility [DC]
Mbps	million bits per second
MHPCC	Maui High Performance Computing Center [DC]
MSRC	Major Shared Resource Center
NAVO	Naval Oceanographic Office [MSRC]
NAWC-AD	Naval Air Warfare Center - Aircraft Division [DC]
NAWC-WD	Naval Air Warfare Center Weapons Division [DC]
NCSA	National Center for Supercomputing Applications
NRL	Naval Research Laboratory [DC]
NRL-DC	Naval Research Laboratory – District of Columbia
OSC	Ohio State Consortium
P-DEAP	Post-Deployment Evaluation and Assessment Process
PET	Programming Environment and Training
PL	Performance Level
POM	Program Objective Memorandum
PRxxxx	Program Review (year)
QoS	quality of service
R&D	research and development
RDT&E	research, development, test, and evaluation

RTTC	Redstone Technical Test Center [DC]
S/AAA	Service/Agency Approval Authority
S&T	Science and Technology
SDP	Service Delivery Point
SDSC	San Diego Supercomputer Center
SIP	signal/image processing [CTA]
SMDC	Space and Missile Defense Command [DC]
SPAWAR	Space and Naval Warfare Systems Command [DC]
SRCAP	Shared Resource Centers Advisory Panel
SSCSD	Space and Naval Warfare Systems Command Center
T&E	Test and Evaluation
TAP	Technical Advisory Panel
TARDEC	Tank-Automotive Research, Development and Engineering Center [DC]
WAN	wide area network
WSMR	White Sands Missile Range [DC]

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	June 2000	Final

4. TITLE AND SUBTITLE
High Performance Computing Modernization Program: Program Review —PR 2000

5. FUNDING NUMBERS
DASW01-98-C-0067
Task Order AK-5-1877

6. AUTHOR(S)
Alfred E. Brenner

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Institute for Defense Analyses (IDA)
1801 N. Beauregard St.
Alexandria, VA 22311-1772

8. PERFORMING ORGANIZATION REPORT NUMBER
IDA Document D-2486

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
High Performance Computing Modernization Program Office
1010 North Glebe Road, Suite 510
Arlington, VA 22201-4795

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, unlimited distribution: 06 November 2000.

12b. DISTRIBUTION CODE
2A

13. ABSTRACT (Maximum 200 words)

The High Performance Computing Modernization Program (HPCMP) was established to advance high performance computing (HPC) for the DoD Science and Technology (S&T) and Test and Evaluation (T&E) communities. In order to ensure the program meets its objectives, the High Performance Computing Modernization Office initiated a formal in-depth review of the HPCMP. This report highlights important issues raised in that review.

14. SUBJECT TERMS
High Performance Computing, Information Technology, Networks, Major Shared Resource Centers, Distributed Centers.

15. NUMBER OF PAGES
58

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102